arXiv:1408.4156v1 [cs.DS] 18 Aug 2014

Efficient Online Strategies for Renting Servers in the Cloud

Shahin Kamali, Alejandro López-Ortiz

University of Waterloo, Canada.

Abstract. In Cloud systems, we often deal with jobs that arrive and depart in an online manner. Upon its arrival, a job should be assigned to a server. Each job has a size which defines the amount of resources that it needs. Servers have uniform capacity and, at all times, the total size of jobs assigned to a server should not exceed the capacity. This setting is closely related to the classic bin packing problem. The difference is that, in bin packing, the objective is to minimize the total number of used servers. In the Cloud, however, the charge for each server is proportional to the length of the time interval it is rented for, and the goal is to minimize the cost involved in renting all used servers. Recently, certain bin packing strategies were considered for renting servers in the Cloud [Li et al. SPAA'14]. There, it is proved that all Any-Fit bin packing strategy has a competitive ratio of at least μ, where μ is the max/min interval length ratio of jobs. It is also shown that First Fit has a competitive ratio of $2\mu + 13$ while Best Fit is not competitive at all. We observe that the lower bound of μ extends to all online algorithms. We also prove that, surprisingly, Next Fit algorithm has competitive ratio of at most $2\mu + 1$. We also show that a variant of Next Fit achieves a competitive ratio of $K \times max\{1, \mu/(K-1)\} + 1$, where K is a parameter of the algorithm. In particular, if the value of μ is known, the algorithm has a competitive ratio of $\mu + 2$; this improves upon the existing upper bound of $\mu + 8$. Finally, we introduce a simple algorithm called Move To Front (MTF) which has a competitive ratio of at most $6\mu + 7$ and also promising average-case performance. We experimentally study the average-case performance of different algorithms and observe that the typical behaviour of MTF is distinctively better than other algorithms.

1 Introduction

Bin packing is a classic problem in the context of online computation. The input is a sequence of *items* of different *sizes* which appear in a sequential, online manner. The goal is to *place* these items into a minimum number of *bins* of uniform capacity so that the total size of items in each bin is no more than the uniform capacity of the bins. It is often assumed that bins have size 1 and items have a positive size no more than 1. The problem is online in the sense that, upon receiving an item, an algorithm should place it into a bin without any knowledge about the (size of) incoming items. A simple online strategy is Next Fit (NF) in which there is a single *open* bin at each time. If an incoming item fits in the open bin, the algorithm places it there; otherwise, it *closes* the open bin and opens a new bin for the item. Clearly, if we want to minimize the number of bins, there is no benefit in closing a bin. First Fit is an online algorithm that never closes a

bin and places an incoming item in the first bin that has enough space for the item; if such a bin does not exist, it opens a new bin. In First Fit, the bins are maintained in the order that they are opened. Best Fit works similarly to First Fit except that it maintains bins in the decreasing order of their fill *level*. The level of a bin is the total size of items placed in the bin. Note that First Fit and Best Fit are greedy algorithm in the sense that they avoid opening new bins unless they have to. The algorithms with this property are called Any Fit algorithm.

In many cloud systems, a set of *jobs* appear in an online manner that should be assigned to servers. Each job has a *load* which defines the amount of resources that it needs. Depending on the application, the load of a job might be defined through its memory requirement, GPU resource usage, bandwidth usage, or a function of all of them. In cloud gaming systems, different instances of computer games are created in an online fashion and run in cloud servers while players interact with the servers via thin clients [8,10]. Here, an instance of a game is a job which, depending on the game and the number of users involved in it, has a load. In case of computer games, the load of a job is mainly defined through the amount of GPU resources that it demands [10].

With the above definition, any online bin packing algorithm can be used to assign jobs to servers. A job of load x can be treated as an item of size x which is assigned to a server (bin) of certain capacity. In this paper, we interchangeably use the terms 'job' and 'item' as well as 'server' and 'bin'. There are, however, distinctions between assigning jobs to servers and the bin packing problem. First, jobs depart the system after they complete; however, the classic bin packing is static in the sense that items are assumed to remain in the bins. When a job arrives, it is not clear when it completes and an algorithm should place it without any knowledge about its departure time. A more important difference between the two problems is that, in the bin packing problem, the objective is to minimize the number of used bins. In other words, we can think of bins as servers that one can *buy* and we would like to minimize the cost by buying a smaller number of servers. In the cloud, however, we want to *rent* the servers from cloud service providers. For example, gaming companies such as OnLive [3] and GaiKai [2] offer cloud gaming services which are execute in public clouds like Amazon EC2 [1]. A rented server is charged by its usage (often in hourly or monthly basis). So, in order to minimize the cost, we need to minimize the total time that servers are rented. In doing so, an algorithm *releases* a server when all the assigned jobs to it are complete.

Definition 1. *In the* server renting problem*, a set of jobs (items) appear in an online manner. Each job has a load (size) that defines the amount of resources that it needs. Upon its arrival, a job should be assigned to a server (bin). Servers have uniform capacity and the total load of jobs assigned to a server should not exceed the capacity. Besides the arrival time, each job has a departure time that indicates when it leaves the system. The length of the interval between the arrival and departure time of a job is referred to as the* length *of the job. Upon the arrival of a job, its length is unknown to the online algorithm. To assign a job to a server, an online algorithm might open (rent) a new server or place it to any of the previously opened servers. When all jobs assigned to a server depart, that server is released. The goal is to minimize the total usage time of servers . More precisely, assuming that an algorithm opens m bins $B_1, \ldots, B_m$, the total cost of the algorithm is $\sum_{i=1}^{m} t_i$, where t_i is the length of time period that B_i has*

been open for. Without loss of generality, we assume the capacity of servers to be 1 and jobs have size at most 1. Also, we assume the length of jobs to be at least Δ and at most $\mu\Delta$ where $\mu \geq 1$.

When studying the server renting problem, we are mostly interested algorithms which have good worst-case and average-case performance. For measuring the worst-case performance, we compare an online algorithm with an optimal offline algorithm OPT that knows the entire sequence (all arrival times, lengths and sizes) in advance. An algorithm is said to be c-competitive (more precisely, *asymptotic* c-competitive) if the cost of serving any input sequence never exceeds c times the cost of OPT within an additive constant.

1.1 Previous Work and Contribution

The Bin Packing problem has been widely studied over the past few decades. It is known that Next Fit is 2-competitive while Best Fit and First Fit are both 1.7-competitive [9]. Generally, any Any Fit algorithm that avoids placing items in the bin with the lowest level is 1.7-competitive (these algorithms are called Almost Any Fit). The Harmonic family of algorithms is another class of bin packing algorithms which are based on placing items of similar sizes together in the same bins. These algorithm generally have better competitive ratios than Any Fit algorithms. The best member of this family is Harmonic++ with a competitive ratio of 1.5888 [11]. However, these algorithms are rarely used in practice since of their poor average-case performance. It is known that no online algorithm can be better than 1.54037-competitive [4].

Coffman et al. [6] studied a dynamic version of the bin packing problem in which items arrive and depart the system. In that variant, it is assumed that the length of any item is revealed upon its arrival. It is proved that the competitive ratio of First Fit is between 2.75 and 2.89 while no online algorithm can do better than 2.5 [6]. For the discrete version of the problem, where each item has size $1/k$ for some integer k, the competitive ratio of Any Fit algorithms is 3 while no online algorithm can do better than 2.48 [5]. Note that in all these results, the objective is to minimize the number of opened bins.

The online server renting problem as defined above was recently introduced by Le et al. [10] (some terms and notations in this paper are also borrowed from [10]). There, the authors prove that no Any Fit algorithm can be better than μ competitive. Recall that μ is the ratio between the length of the largest and the smallest item in the sequence. They also proved that the competitive ratio of First Fit is $\frac{k}{k-1}\mu + \frac{6k}{k-1} + 1$ when the size of items are upper bounded by $1/k$ (k can be any positive value). In particular, when $k = 1$ (when there is no restriction on item sizes), First Fit is $2\mu + 13$-competitive. On the other hand, they prove that Best Fit is not competitive, i.e., it does not achieve a constant competitive ratio. This result is somewhat surprising as Best Fit is usually considered to be the superior algorithm for many bin packing applications. In [10] a slight modification of the First Fit algorithm is introduced which achieves a competitive ratio of $\frac{8}{7}\mu + 55/7$ when the value of μ is not known to the algorithm and a competitive ratio of $\mu + 8$ when the value of μ is known.

In this paper, we study the server renting problem and improve some of the results presented in [10]. We first observe that the lower bound of μ presented for competitiveness of any Any Fit algorithm can be extended to any online algorithm. The focus of [10] has been on studying Any Fit algorithms. In the standard bin packing, Any Fit algorithms have an advantage over bounded-space algorithms which close the bins (since there is no harm in keeping bins open). However, we show that this is not the case for the server renting problem and it makes sense to close some servers to avoid placing new items in the bins opened by older items. In particular, we show that the competitive ratio of the Next Fit algorithm is at most $\frac{k}{k-1}\mu+1$ when items are smaller than $1/k$ and $2\mu+1$ in the general case. Note that this is much better than the ratio $2\mu+13$ of First Fit. We also introduce variants of Next Fit which achieve a ratio of $K \times \max\{1, \frac{\mu}{K-1}\}+1$, where K is a parameter of the algorithm and can be any positive value. In particular, if the value of μ is known, we get an algorithm with competitive ratio of $\mu + 2$ which is better than $\mu + 8$ of the algorithm presented in [10].

Although Next Fit has a superior competitive ratio compared to Best Fit and First Fit, unfortunately, it has a poor average-case performance relative to these algorithms. Our experiments indicate that, for sequences generated uniformly at random, Best Fit performs generally better than the other two algorithms. To address this issue, we introduce a simple Any Fit algorithm called Move To Front (MTF) which outperforms Best Fit and other algorithms on random sequences. Moreover, in contrast to Best Fit, MTF is competitive and has a competitive ratio of at most $6\mu + 7$.

2 Preliminaries

In this section, we present some basic results about the server renting problem. First, we show that any online algorithm for the server renting problem has a competitive ratio of at least μ. Our lower bound sequence is similar to that of [10] and is composed of items with uniform sizes and different lengths.

Theorem 1. *The competitive ratio of any online algorithm for the server renting problem is at least* $\frac{\mu}{1+\epsilon(\mu-1)}$ *where* ϵ *is a lower bound for the size of items.*

Proof. Recall that the lengths of all items are at least Δ and at most $\mu\Delta$. Consider a sequence which is defined through phases. Each phase starts with $\frac{1}{\epsilon^2}$ items of size ϵ. To place these items, any algorithm has to open at least $1/\epsilon$ bins. At time Δ, $\frac{1}{\epsilon^2} - \frac{1}{\epsilon}$ items depart in an adversarial manner so that there is a single of item of size ϵ in $1/\epsilon$ bins (some bins might get released at this time). The remaining items stay for a period of length $\mu\Delta$ and the online algorithm keeps a single bin for each of them. At time $\mu\Delta$, all items depart and the phase ends. The cost of the online algorithm for each phase is at least $\mu\Delta/\epsilon$ (it keeps $1/\epsilon$ bins for a period of $\mu\Delta$). OPT places items which have length $\mu\Delta$ together in a single bin and incurs a cost of $\mu\Delta$ for them. Other $\frac{1}{\epsilon^2} - \frac{1}{\epsilon}$ items are placed tightly together in $1/\epsilon - 1$ bins for a period of length Δ (after which all they leave and their bins get closed. The cost of OPT for these items will be $\Delta/\epsilon - \Delta$. In total, the cost of OPT will be $\mu\Delta + \Delta/\epsilon - \Delta$ and the competitive ratio of the algorithm will be $\frac{\mu\Delta/\epsilon}{\mu\Delta+\Delta/\epsilon-\Delta} = \frac{\mu}{1+\epsilon(\mu-1)}$. □

Next, we introduce two lower bounds for the cost of OPT for serving any input sequence. We say an item x is *active* at time t if t lies in the interval between the arrival and the departure time of a. Let the *span* of an input sequence σ denote the length of time in which at least one item in σ is active. Clearly, the cost of any algorithm (including the offline OPT) for serving σ is at least equal to the span of σ. Define the *resource utilization* of an item as the product of its size and its length. This way, the cost of any algorithm for σ is at least equal to $util(\sigma)$, that is, the total resource utilization of items in σ. So, the cost of an optimal algorithm for an input sequence is always lower bounded by the span of the sequence and also by the total utilization of the sequence.

Proposition 1. *For any input sequence σ, the cost of an optimal offline algorithm* OPT *is at least equal to $span(\sigma)$ and $util(\sigma)$, namely, the span of σ and also the total resource utilization of items in σ.*

When we allow arbitrary small items, Theorem 1 indicates that all algorithms have a competitive ratio of at least μ. This suggests that when item sizes are larger than a fixed value, better competitive ratios can be achieved. Consider a sequence σ in which all item sizes are larger than $1/k$ for some positive value k. The cost of any algorithm is at most equal to the total length of all items denoted by $L(\sigma)$ (which happens when no two items share a bin). On the other hand, the total resource utilization of items, and consequently cost of OPT, is at least $L(\sigma)/k$. So, we get the following.

Proposition 2. *[10] When items sizes are lower bounded by $1/k$ (k is a positive value), the competitive ratio of any online algorithm for the server renting problem is at most equal to k.*

3 Next Fit Algorithm

In this section, we analyze the Next Fit algorithm for the server renting problem. Recall that for the bin packing problem Next Fit keeps exactly one bin open at any given time. If an incoming item does not fit in the open bin, it closes the bin and opens a new bin. For the server renting problem, we distinguish between *closing* and *releasing* a bin. When an item does not fit in the open bin, the algorithm closes the bin and does not refer to it. Such a bin remains in the system (i.e., is being rented) until all items which are placed there depart and it becomes released.

Example 1. Consider sequence $\langle a = (0.3, 1, 5), b = (0.4, 2, 6), c = (0.4, 3, 5), \ldots \rangle$ of items. The first element of each tuple indicates the size and the second and third respectively indicate arrival and departure times of an item. At time 1, item a arrives and is placed in the single open bin. At time 2, item b arrives and is placed in the same bin (the level of the bin will be 0.7). At time 3, item c arrives which does not fit in the open bin; hence, the current open bin is closed and a new bin is opened for c. The closed bin remains in the system (and a rental cost is paid for it) until time 6 where item b departs and the bin gets released.

Theorem 2. *The competitive ratio of Next Fit for the server renting problem is at most $\frac{\mu}{1-1/k} + 1$ for serving sequences in which item sizes are no more than $1/k$. If $k < 2$, the ratio is at most $2\mu + 1$.*

Proof. Consider an arbitrary sequence σ and assume Next Fit opens m bins $B_1, \ldots, B_m$ for serving σ. Let st_i denote the length of the time interval at which the server B_i has been rented, i.e., the time after it is opened and before it is released. We refer to this period as *stretch* of B_i. Let $\text{NF}(\sigma)$ denote the cost of Next Fit for serving σ; we have $\text{NF}(\sigma) = \sum_1^m st_i$. The stretch of B_i can be partitioned into two period. First, the interval between its opening time and when Next Fit closes the bin. The second period is the time between the bin gets closed and when it gets released. Let st_i^1 and st_i^2 denote the lengths of first and second period of B_i ($st_i^1 + st_i^2 = st_i$). If a bin gets closed before being released, the second period will be empty ($st_i^2 = 0$) (see Figure 1).

Let $p \leq m$ denote the number of bins which are closed before being revealed. We call these bins *critical* bins and for them we have $st_i^2 \neq 0$. The main observation is that, when a bin gets closed, it takes a time of length at most μ before it gets released, i.e., the second period of each bin has a length of at most μ ($st_i^2 \leq \mu$ for all i). This is because no new item is placed in the bin in the second period. So, the total rental time for the second period of all bins is no more than $p \times \mu$. On the other hand, the total rental time of the first periods of all bins is no more than the span of input sequence. This is because the first period of a bin starts when that of previous bin is finished (no two bins are in their first stretch period at the same time). So we have

$$\text{NF}(\sigma) = \sum_{i=1}^{m} st_i = \sum_{i=1}^{m} st_i^1 + \sum_{i=1}^{p} st_i^2 \leq span(\sigma) + p \times \mu\Delta \tag{1}$$

Assume that all items in σ are smaller or equal to $1/k$ for $k \geq 2$. At the time of being closed, all critical bins have a level of at least $1 - 1/k$ (otherwise the item that caused opening of a new bin could fit in such a bin). This implies that the number of critical bins (i.e., p) cannot be more than $\frac{\omega(\sigma)}{1-1/k}$ where $\omega(\sigma)$ is the total size of items in σ. Let $util(\sigma)$ denote the total resource utilization of items in σ. Since the length of each item is at least Δ, we have $util(\sigma) \geq \omega(\sigma) \times \Delta$, and by Proposition 1 $\omega(\sigma) \leq \text{OPT}(\sigma)/\Delta$. Consequently, $p \leq \frac{\text{OPT}(\sigma)}{(1-1/k)\Delta}$. Also, by Proposition 1, $span(\sigma) \leq \text{OPT}(\sigma)$. Plugging these into Equation 1, we get the following inequality which completes the proof.

$$\text{NF}(\sigma) \leq \text{OPT}(\sigma) + \frac{\text{OPT}(\sigma)}{1 - 1/k} \times \mu$$

Next, assume $k \leq 2$. We define the *amortized level* of a critical bin B as the sum of the size of the item that closes B and the total size of items in B (at the time that Next Fit closes it). By definition of NF, the amortized level of all critical bins is more than 1. At the same time, the size of each item is added at most twice in the total amortized cost (once as a part of a critical bin and once as the time that closes a critical bin). Hence, the total sum of the amortized levels of all critical bins is at most twice the total size of sequence. This implies that the number of critical bins is no more than twice the total size of items in σ, i.e., $p \leq 2\omega(\sigma) \leq 2\,\text{OPT}(\sigma)/\Delta$. Applying this into Equation 1, we get the following inequality which completes the proof.

$$NF(\sigma) \leq \text{OPT}(\sigma) + 2\,\text{OPT}(\sigma) \times \mu$$

□

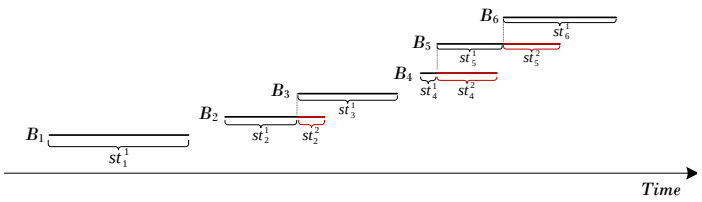

Fig. 1. The stretch of bins in a packing of the Next Fit algorithm. In this example, bins B_2, B_4, and B_5 are critical bins. The second periods of bins are highlighted in red. Note that the total lengths of the first periods defines the span of the input sequence.

3.1 Improving the Competitive Ratio: Modified Next Fit Algorithm

In this section, we modify the Next Fit algorithm to improve its competitive ratio. Intuitively speaking, the competitive ratio of Next Fit improves for sequences formed by small items. On the other hand, as Proposition 2 implies, when all items are large, the competitive ratio is independent of μ. This suggest that the competitive ratio might get improved when large and small items are treated separately. A similar approach is used in [10] to improve the competitive ratio of First Fit.

Modified Next Fit: The algorithm has a parameter $K \geq 2$ and treats items smaller than $1/K$ using the Next Fit strategy. Items larger or equal to $1/K$ are treated separately also using the Next Fit strategy.

Theorem 3. *The competitive ratio of Modified Next Fit with parameter K is at most $K \times max\{1, \mu/(K-1)\} + 1$.*

Proof. Consider a sequence σ and let σ^s and σ^l denote the subsequences of σ respectively formed by items smaller and larger or equal to K. Recall that the resource utilization of an item is the product of its length and its size, and the total resource utilization of all items in a sequence is a lower bound for the cost of OPT for that sequence. As Proposition 2 suggests, the number of opened bin by Modified Next Fit for items in σ^l is no more than $k \times util(\sigma^l)$, where $util(\sigma^l)$ is the total utilization of items in σ^l.

For placing σ^s, as the proof of Theorem 2 suggests, the algorithm incurs a cost of at most $\mu\Delta \times \frac{\omega(\sigma^s)}{1-1/K} + span(\sigma^s)$, where $\omega(\sigma^s)$ is the total size of item in σ^s. This will be no more than $\mu\Delta \times \frac{util(\sigma^s)}{\Delta(1-1/K)} + span(\sigma^s)$, where $util(\sigma^s)$ is the total resource utilization of σ^s (this is because the length of all items is at least Δ). In total, the cost of the algorithm will be at most $K \times util(\sigma^l) + \mu \times \frac{util(\sigma^s)}{1-1/K} + span(\sigma^s)$. This is no more than $K \times util(\sigma) \times \max\{1, \frac{\mu}{K-1}\} + span(\sigma)$ where $util(\sigma)$ is the total utilization of items in σ. Since $util(\sigma)$ and $span(\sigma)$ are lower bounds for the cost of OPT, we can conclude the cost of Modified Next Fit is at most $K \times \max\{1, \frac{\mu}{K-1}\} + 1$. □

When the value of μ is known to the algorithm, we can define K to be $\mu + 1$. In this case, the competitive ratio of Modified First Fit will be $\mu + 2$.

Proposition 3. *When the value of μ is known, there is an online algorithm with competitive ratio of $\mu + 2$.*

4 Toward Practical Algorithms: Move To Front Algorithm

In the previous sections, we showed that Next Fit and a variant of that have promising competitive ratios. These kinds of worst-case guarantees are important in theoretical analysis of the problem. Nevertheless, in practice, beside worst-case guarantees, we are interested in algorithms that also have promising average-case performance. For example, in the case of the classic bin packing problem, Best Fit and First Fit are preferred over other algorithms in most applications because they have acceptable worst-case performance (although not as good as the Harmonic family of algorithms) and superior average-case performance. We examined different packing algorithms to evaluate their average-case performance. Our experiments are presented in Section 5 and show that, on average, Best Fit has an evident advantage over First Fit, and First Fit is better than Next Fit. These results are in contrast with competitive results (recall that Best Fit is not competitive at all) and indicate that the worst-case behaviour and average-case behaviour of these algorithms are quite different.

In this section, we introduce and evaluate the Move-To-Front (MTF) algorithm for the server renting problem. We prove that, unlike Best Fit, this algorithm is competitive (i.e., provides a worst-case guarantee). Our experiments indicate that MTF performs better than all other algorithms on randomly generated sequences. MTF is a simple Any Fit algorithm and runs as fast as BF and FF. Hence, we believe that this is the best algorithm for the server renting problem.

> *Move To Front:* The algorithm maintains a list of open bins. When a new item x arrives, the algorithms checks the bins one by one, starting from the front of the list, until it finds a bin that has enough space for x. If no bin has enough space, a new bin is opened for x. After placing x in a bin, that bin is moved to the front of the list.

MTF tends to place items which arrive almost at the same time in the same bins. By avoiding placing old an new items in the same bins, the algorithm avoids bad situation in which a bin is open for a small item x, and just before x departs another small item y is placed in the bin.

Theorem 4. *Move To Front has a competitive ratio of at most $6\mu + 7$.*

Proof. Consider the final packing of the algorithm for a sequence σ. We assume that the sequence is *continuous* in the sense that it has a continuous span, i.e., at each time there is at least one active item and any algorithm maintains at least one open bin. For sequences which are not continuous, at some point, all bins of MTF and OPT are closed. In this case, we can divide the sequence into continuous subsequences and apply the same argument for each of them.

We divide the span of the sequence into *periods* of length $(\mu + 1)\Delta$ (except the last period which might be shorter). For each bin B, we define *head*, *tail*, and *body* of B as

follows. If B is opened and closed in the same period, its head is its stretch (interval between its opening and closing) while its body and tail are empty intervals. Otherwise, head of B is the interval between when B and the end of the period in which it is opened. Similarly, tail of B is the interval between the start of the period in which it is closed and when B is closed. Body of B is the interval between its head and tail. Figure 2 provides an illustration. Let $head(B)$, $body(B)$, and $tail(B)$ indicate the lengths of respectively head, body, and tail of B. For the cost that MTF incurs for B (stretch of B) we have:

$$stretch(B) = head(B) + body(B) + tail(B) \leq 2(\mu + 1)\Delta + body(B)$$

Assume there are m bins opened by MTF. The algorithm incurs a cost of at most $2 \times (\mu + 1)$ for head and tail of each bin. We will have:

$$\text{MTF}(\sigma) \leq 2m(\mu + 1)\Delta + \sum_{b=1}^{m} body(B_b)$$

Assume there are q periods in the packing. For each period P_i ($1 \leq i \leq q$), let $\alpha(P_i)$ denote the number of bins which have their body in P_i, that is, bins which are open at the beginning of the period and remain open till the end. Note that $\alpha(P_i) \geq 1$. MTF incurs a cost of $\alpha(P) \times (\mu + 1)$ for body of all bins in P. We will have:

$$\text{MTF}(\sigma) \leq 2m(\mu + 1)\Delta + (\mu + 1)\Delta \sum_{i=1}^{q} \alpha(P_i) \tag{2}$$

Next, we consider the cost of OPT for packing σ. We prove the following claims:

Claim 1: For the number of bins opened by MTF we have $m \leq 2\,\text{OPT}(\sigma)/\Delta + 1$.
Claim 2: For each period P, if $\alpha(P) = 1$, OPT incurs a cost of at least $(\mu + 1)\Delta$.
Claim 3: For each period P, if $\alpha(P) \geq 2$, OPT incurs a cost of at least $(\alpha(P) - 1)\Delta/2$.

Claim 1 implies that the first term in Equation 2 is upper bounded by $4(\mu+1)\,\text{OPT}(\sigma)+ 2(\mu+1)\Delta$. Claims 2 implies that in the specified periods, MTF and OPT incur the same costs. Claim 3 implies that $\alpha \leq 2\,\text{OPT}(P)/\Delta + 1$ where $\text{OPT}(P)$ is the cost inured by OPT in period P. Consequently, the second term in Equation 2 is upper bounded by

$$\begin{aligned}(\mu + 1)\Delta \sum_{i=1}^{q} (2\,\text{OPT}(P_i)/\Delta + 1) &= 2(\mu + 1) \sum_{i=1}^{q} \text{OPT}(P_i) + q(\mu + 1)\Delta \\ &< 2(\mu + 1)\,\text{OPT}(\sigma) + span(\sigma) + (\mu + 1)\Delta\end{aligned}$$

The last inequality holds because we have divided the stretch of σ into q periods with equal length of $(\mu + 1)\Delta$ (except the last period which might be shorter). Adding both terms in Equation 2, we get:

$$\begin{aligned}\text{MTF}(\sigma) &\leq 4(\mu + 1)\,\text{OPT}(\sigma) + 2(\mu + 1)\Delta + 2(\mu + 1)\,\text{OPT}(\sigma) + \text{OPT}(\sigma) + (\mu + 1)\Delta \\ &= (6\mu + 7)\,\text{OPT}(\sigma) + 3(\mu + 1)\Delta\end{aligned}$$

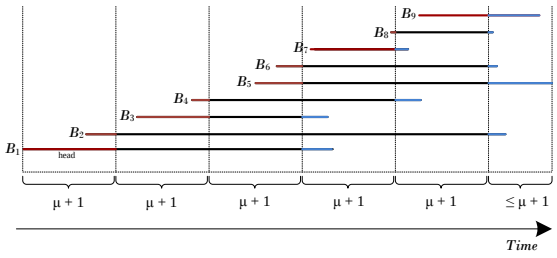

Fig. 2. The span of an input sequence is divided into periods. The red and blue intervals respectively indicate heads and tails of bins while black intervals are the bodies of the bins.

This implies that the competitive ratio of the algorithm is at most $6\mu + 7$ (note that $3(\mu + 1)\Delta$ is a constant). To complete the proof, it remains to show the above claims hold.

For Claim 1, consider the ordering in which bins are opened and define the amortized weight of a bin (except the last bin in the ordering) as the total size of items in the bin plus the item that causes opening of the next bin. With this definition, every bin (except the last bin) has amortized weight of more than 1. Thus, the total amortized cost of all bins will be more than $m - 1$. Each item is counted at most twice in the total amortized cost (once as the item that opens a new bin and once as the member of a bin which cannot fit a new item). Assuming $\omega(\sigma)$ is the total size of items, we will have $m - 1 < 2\omega(\sigma)$. Note that $\omega(\sigma) \leq util(\sigma)/\Delta \leq \textsc{Opt}(\sigma)/\Delta$. Hence, we get $m \leq 2\,\textsc{Opt}(\sigma)/\Delta + 1$.

For Claim 2, note that at each time, OPT maintains at least one open bin; otherwise, the sequence is not continuous.

For Claim 3, let t denote the start time of P and let B^* denote the set of the $\alpha(P)$ bins which have their body stretched along P. Consider the time interval $[t+\Delta, t+(\mu+1)\Delta)$. In this interval, any of the bins in B^* receive at least one new item; otherwise, the algorithm would have closed the bin (recall that the length of any item is at most $\mu\Delta$). For each bin B in B^*, except the last bin in the list maintained by the algorithm right before time $t + \Delta$, let t_B indicate the time that the bin B' receives an item for the first time (in the interval $[t + \Delta, t + (\mu + 1)\Delta)$). Here, B' is the bin that is placed right after B in the mentioned ordering. Define the *critical set* of B as the set of items in B at time t_B plus the item that was placed in B'. Note that the total size of items in the critical set of each bin is more than 1. Hence, the critical items of each bin have a resource utilization of more than Δ in the interval $[t, t + (\mu + 1)\Delta)$. Since each item belongs to critical sets of at most two bins, the total resource utilization of critical items is at least $(\alpha(P) - 1)\Delta/2$ in the interval $[t, t + 1 + \mu)$. Note that the resource utilization is a lower bound for the cost of OPT in the same interval. □

In fact, the above proof can be extended to any algorithm that maintains a list of bin and places an incoming item in the first bin which has enough space. Such an algorithm might update the list after *placing* items (as MTF does). In particular, the above analysis also applies for the First Fit algorithm. For the Best Fit algorithm, the above analysis

fails because the order of bins changes when items *depart*. Recall that Best Fit is not competitive at all. The bad sequences that cause unbounded competitive ratio of Best Fit cause the algorithm keep multiple bins open for an arbitrary long time with an arbitrary small level of ϵ. This cannot happen for MTF and First Fit.

5 Average-case Analysis: An Experimental Study

In this section, we study the performance of MTF compared to other algorithms for the server renting problem on randomly-generated sequences. We discretize the problem by assuming that the size of bins is an integer E and items have integer sizes in the range $[1, E]$. Moreover, we assume items arrive in discrete time-steps in the range $[1, T - \mu]$ and their length is in the interval $[1, \mu]$. Here, T denotes the span of generated sequences. We examine different values of μ and T for sequences of fixed length. This way, T is a measure of *sparsness* and defines the rate at whcih the items arrive. Table 1 gives details of the datasets that we generated for our experiments. In all cases, both size and length of items are randomly and independently taken from the indicated ranges (assuming a uniform distribution). For each setting of the problem, we run different algorithms on 10^3 randomly generated sequences. For each sequence, we compute the resource utilization of the sequence as a lower bound for the cost of OPT (see Proposition 1). We use the ratio between the cost of an algorithm and the resource utilization as a measure of performance.

The algorithms that we considered in the experiments are Next Fit, Modified Next Fit, First Fit, Modified First Fit, Harmonic, Best Fit, and Move To Front. We define the parameters of Modified Next Fit and Modified First Fit to be respectively $E/(\mu + 1)$ and $E/(\mu + 7)$. These values ensure that these algorithms achieve their best possible competitive ratio (see Section 3.1 and [10]). Note that the value of μ is not known to the online algorithm and these algorithms are semi-online in this sense. We also consider the Harmonic algorithm which classify items by their sizes (using harmonic intervals) into K classes and applies the Next Fit strategy for placing items of each class; we assume $K = 10$ in our experiments. A straightforward analysis shows that the competitive ratio of the Harmonic algorithm is as good as Modified Next Fit (with the same parameter K). However, similar to the bin packing problem, for the server renting problem, Harmonic seems to have a poor average-case performance.

Figure 3 shows the average-case performance ratio among all sequences for different algorithms. In most cases, Move To Front is the best algorithm. Intuitively, there are two factors which define the quality of a packing. One is how well items are *aligned*

Parameter	Description	Value	Note
n	length of sequences	10^5	Number of items to be packed
μ	maximum length of items	1,2,5,10,100	Lengths are picked from the range $[1, \mu]$
T	span of sequence	$10^3, 10^4, 10^5$	Arrival times are picked from the range $[1, T - \mu]$
E	bin capacity	10^3	Sizes are picked from the range $[1, E]$

Table 1. A summary of the experimental settings.

to each other. Informally speaking, a packing is well-aligned if items that arrive at the same time are placed together; this ensures that, on expectation, all items of a bin depart also at (almost) the same time. Thus, there is a more chance of saving cost through closing bins. The second factor in defining the quality of a bin is how well the items are packed together according their sizes. Clearly, if items are tightly packed together, there is a save in cost by avoiding opening new bins.

By definition, Next Fit results in well-aligned packings; however, it does not packs items as tightly as Any Fit algorithms do. On the other side, Best Fit results tight packings which are not necessarily well-aligned. Move To Front provides a compromise. The packing of MTF are well-aligned because items placed in the most recent bin have almost same arrival time; meanwhile avoiding to place items in other bins give the a chance of being closed (and saving cost thoroughly). At the same time, as an Any Fit algorithm, MTF places items almost tightly and does not open a large number of bins as Next Fit does. For smaller values of μ, it is more important to achieve well-aligned packings. this is because, when items have roughly same length, there is more benefit in placing them together according to their arrival time since it ensures that they depart at almost same time and their hosting bin gets released. This is particularly more evident for sequences with small span ($T = 1,000$). In these sequences, many items appear at the same time and almost all algorithms result in relatively good packing (regarding item sizes). As a results, for smaller values of μ and T, Next Fit performs better relative to other algorithms. In particular, when $\mu = 1$ and $T = 1,000$, it slightly outperforms MTF. For larger values of μ, it is more important to avoid opening new bins; this is because each bin remains open for a relativity long period of time and one should avoid opening a new bin as much as possible. As a result, when μ is large ($\mu = 100$), Best Fit outperforms Move To Front.

	T=1,000					T=10,000					T=100,000				
	μ=1	μ=2	μ=5	μ=10	μ=100	μ=1	μ=2	μ=5	μ=10	μ=100	μ=1	μ=2	μ=5	μ=10	μ=100
Next Fit	**1.4011**	1.4618	1.4970	1.5113	1.5255	1.7186	1.6564	1.5810	1.5473	1.5263	1.9539	1.9291	1.8721	1.8014	1.5496
Modified Next Fit	1.4392	1.4820	1.4780	1.5023	1.5253	1.8069	1.8061	1.7223	1.6383	1.5366	1.9762	1.9738	1.9313	1.8601	1.5631
First Fit	1.5647	1.4561	1.3448	1.2929	1.2255	1.7485	1.6635	1.5288	1.4295	1.2624	1.9544	1.9292	1.8705	1.7950	1.4132
Modified First Fit	1.6666	1.5335	1.3952	1.3233	1.2287	1.8362	1.7622	1.6330	1.5148	1.2768	1.9726	1.9534	1.9042	1.8352	1.4293
Harmonic	1.7598	1.6937	1.5848	1.5143	1.4235	1.9555	1.9270	1.8842	1.8195	1.5170	1.9946	1.9903	1.9797	1.9726	1.8294
Best Fit	1.6659	1.5027	1.3509	1.2637	**1.1151**	1.7401	1.6585	1.5342	1.4359	**1.2107**	1.9540	1.9287	1.8696	1.7935	1.4164
Move To Front	1.4113	**1.3921**	**1.3094**	**1.2560**	1.1612	**1.7134**	**1.6323**	**1.5036**	**1.4110**	1.2251	**1.9536**	**1.9283**	**1.8689**	**1.7913**	**1.4005**

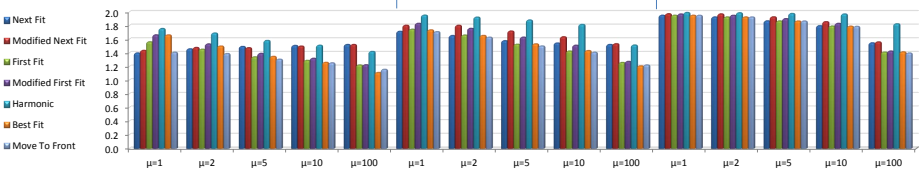

Fig. 3. Average-case performance ratio of major server renting problem algorithms, assuming a uniform distribution for the length and size of items. The bold numbers indicate the best algorithm for different values of T and μ. In all cases MTF is the best or second to the best algorithm. To make comparison easier, the numbers are plotted into a bar diagram.

6 Concluding Remarks

In this paper, we showed that the Next Fit algorithm provides promising worst case guarantees for the server renting problem. We expect that the same holds for other bounded-space algorithm, e.g., Harmonic or BBF_2 of [7]. Unfortunately, these algorithms do not have good average-case performance. To address this issue, we introduced the MTF algorithm which is a simple and fast Any Fit algorithm which can be regarded as an alternative to the Best Fit and First Fit algorithms. Our experiments indicate that MTF outperforms other algorithms on the average case for placing sequences that are generated randomly. The closest counterpart of MTF (regarding average case performance) is the Best Fit algorithm which is not good in the worst case as it is not competitive at all. In contrast to Best Fit, we proved that MTF is competitive and has a competitive ratio of at most $6\mu + 7$. We believe this upper bound is not tight and the competitive ratio of MTF can be improved to $2\mu + 1$ of NF; we leave this as a future work. Another promising direction for future work is to to provide theoretical upper bounds for the average-case performance of MTF on sequences that follow arbitrary distributions.

References

1. Amazon EC2. http://aws.amazon.com/ec2/, accessed: 2014-08-14
2. Gaikai. http://www.gaikai.com/, accessed: 2014-08-14
3. OnLive. http://www.onlive.com/, accessed: 2014-08-14
4. Balogh, J., Békési, J., Galambos, G.: New lower bounds for certain classes of bin packing algorithms. Theoret. Comput. Sci. 440–441, 1–13 (2012)
5. Chan, W.T., Lam, T.W., Wong, P.W.H.: Dynamic bin packing of unit fractions items. Theoret. Comput. Sci. 409(3), 521–529 (2008)
6. Coffman, E.G., Garey, M.R., Johnson, D.S.: Dynamic bin packing. SIAM J. Comput. 12, 227–258 (1983)
7. Csirik, J., Johnson, D.S.: Bounded space on-line bin packing: Best is better than first. pp. 309–319 (1991)
8. Huang, C.Y., Hsu, C.H., Chang, Y.C., Chen, K.T.: Gaminganywhere: An open cloud gaming system. In: Proc. the 4th ACM Multimedia Systems Conference. pp. 36–47. MMSys '13 (2013)
9. Johnson, D.S., Demers, A., Ullman, J.D., Garey, M.R., Graham, R.L.: Worst-case performance bounds for simple one-dimensional packing algorithms. SIAM J. Comput. 3, 256–278 (1974)
10. Li, Y., Tang, X., Cai, W.: On dynamic bin packing for resource allocation in the cloud. pp. 2–11 (2014)
11. Seiden, S.S.: On the online bin packing problem. J. ACM 49, 640–671 (2002)